



- 1 -

ANOMALY NOTIFICATION CONTROL IN DISK ARRAY

INCORPORATION BY REFERENCE

The present application claims priority from Japanese application No. 2004-027490 filed on February 4, 2004 the content of which is hereby incorporated by reference into this application.

BACKGROUND OF THE INVENTION

The present invention relates to a disk array incorporating different kinds of disk drives. More specifically the present invention relates to a disk array which, in the event of a failure of a part of the disk drives, can perform sparing by using different kinds of disks and also to a sparing method.

A disk array accommodates a large number of disk drives. Should a part of these disk drives fail, a normal operation of the disk array cannot be guaranteed. As a means for improving a fault tolerance of the disk array, sparing may be used. The sparing involves preparing spare disk drives in a disk array in advance and, when a failure is detected, quickly disabling the failed disk drive and placing a spare disk drive in operation. After sparing is effected, an anomaly is notified to an administrator to prompt him to perform a maintenance service. By replacing the failed disk drive with a normal spare disk drive in

this manner, the disk array can be maintained without stopping its operation.

JP-A-5-100801 discloses a technique which, when the number of access errors in a disk drive exceeds a predetermined value, disables the disk drive preventively before it fails and swaps it with a spare disk drive. JP-A-2002-297322 discloses a technique which, in the event of a failure, distributively stores data from the disabled disk drive in a plurality of spare disk drives.

SUMMARY OF THE INVENTION

There are a variety of kinds of disk drives with different characteristics, such as fibre channel disk drives with a fibre channel interface (hereinafter referred to as "FC disk drives") and serial disk drives with a serial interface (referred to as "SATA disk drives"). In a disk array, the use of different kinds of disk drives can not only take advantage of features of these disk drives but also compensate for their shortcomings. To perform sparing in such a disk array, it is desired that spare disk drives be prepared for each kind of disk drive.

However, there is a limit on the number of disk drives that can be installed in the disk array. Thus, in preparing spare disk drives for each kind a problem arises that a sufficient number of spare disk drives may not be available for each kind. With

sufficient numbers of spares not available, a failure of even a small quantity of disk drives, which reduces the number of remaining spare disk drives, makes it necessary to perform maintenance service frequently,

5 increasing a maintenance overhead, which should be avoided. Under these circumstances, the present invention enables sparing in a disk array incorporating different kinds of disk drives without causing an excessive increase in a maintenance overhead.

10 The present invention concerns a disk array which has installed in a disk array rack a plurality of disk drives and controllers for controlling data read/write operations to and from the disk drives, with the disk drives and the controllers interconnected with
15 cables. In this disk array there are different kinds of disk drives with different characteristics. With this invention, whether a disk drive is to be disabled or not is decided by the controllers based on the number of errors that occur during the read/write
20 operations in each disk drive. If it is decided that a certain disk drive be disabled, sparing processing is executed to allocate a part of disk drives as a spare for the disk drive that is going to be removed from service. The disk drives used for sparing may or may
25 not be of the same kind as the disk drives to be disabled.

For example, the present invention provides a disk array comprising: a disk array rack; a plurality

of disk drives installed in the disk array rack; a controller installed in the disk array rack to control data reads and writes to and from the disk drives; and cables connecting the controller with the disk drives;

5 wherein the disk drives comprise first disk drives and second disk drives with an interface different from that of the first disk drives; wherein the controller, when it decides that one of the first disk drives fails, performs sparing on the failed first disk drives

10 by using the second disk drives.

As a result of disabling a disk drive, the controller notifies the occurrence of the disabled state to a predetermined notification destination at a predetermined notification timing. In this invention

15 the notification timing is set so that the notification resulting from the sparing performed between the disk drives of different kinds is issued earlier than the notification resulting from the sparing performed between the disk drives of the same kind. As an

20 example, the anomaly notification may be issued immediately when the sparing is done between different kinds of disk drives but may be delayed a certain period of time when the sparing is done between the same kinds.

25 With this invention, by permitting sparing between different kinds of disk drives, it is possible to secure a sufficient number of disk drives that can be used as spares and thereby avoid the maintenance

interval becoming short. However, the sparing between different kinds of disk drives may not be able to secure a sufficient performance due to a characteristic difference between these disk drives. Taking this 5 problem into account, this invention advances the notification timing for the sparing between different kinds of disk drives to minimize performance reduction of the disk array.

In this invention it is preferred that the 10 execution of the sparing between disk drives of the same kind be given priority over the execution of the sparing between different kinds. This can minimize a performance reduction of the disk array caused by sparing.

15 In this invention, the notification timing may be set based on at least the number of disabled disk drives or the number of disk drives available for the sparing. For instance, when the number of disabled disk drives exceeds a predetermined value or when the 20 number of spares falls below a predetermined value, the anomaly notification may be issued. This eliminates a possibility of bringing about a situation in which the disk array is forced to be shut down because of unduly delayed notification.

25 In this invention, other failures than the disabled state in the disk array may be notified. In that case, if a failure other than the disabled state should occur before the notification timing is reached, this

failure may be notified along with the disabled state. This allows maintenance on a variety of failures to be performed at the same period, reducing the maintenance burden.

5 In this invention, when performing sparing between different kinds of disk drives, the allocation of disk drives may be controlled so as to compensate for a characteristic difference between different kinds of disk drives. In the case of sparing between FC disk
10 drives and SATA disk drives, for example, a failed FC disk drive may be subjected to sparing by parallelly assigning a plurality of SATA disk drives. Parallel assignment means an arrangement that allows parallel accesses to the plurality of disk drives. Generally,
15 SATA disk drives have a slower access speed than FC disk drives. The parallel allocation therefore can prevent a reduction in access speed.

Conversely, when a serial disk drive is disabled, a plurality of fibre channel disk drives may
20 be serially assigned. Generally, FC disk drives have a smaller capacity than SATA disk drives. By serially assigning the FC disk drives, it is possible to minimize a capacity reduction as a result of sparing.

This invention can be applied to a variety of
25 disk arrays, including one which incorporates a combination of FC disk drives and SATA disk drives. In this configuration, it is preferred that the disk array have a converter to convert a serial interface of each

SATA disk drive into a fibre channel interface. This arrangement can transform the interfaces of various disk drives into a unified interface, i.e., the fibre channel.

5 Further, dual paths may be employed to improve a fault tolerance of the disk array. That is, a plurality of fibre channels may be formed by providing a plurality of controllers, interconnecting the controllers through fibre channel cables, and
10 connecting each of the controllers with individual disk drives through the fibre channel cables. As to the SATA disk drives, dual paths can be formed by providing a selector which selects a connection destination of the SATA disk drives among a plurality of fibre channel
15 loops.

This invention can be implemented not only as a disk array but also as an anomaly notification control method in a disk array. For example, an anomaly notification control method for controlling a
20 notification of an anomaly that has occurred in a disk array may comprise: a disk array rack; a plurality of disk drives installed in the disk array rack; and a controller installed in the disk array rack to control data reads and writes to and from the disk drives;
25 wherein the disk drives comprise a plurality of kinds of disk drives with different characteristics; wherein the controller executes: a decision step of evaluating errors that occur during reads and writes to and from

each of the disk drives and deciding whether each disk drive needs to be disabled or not; a sparing control step of controlling sparing processing which, when it is decided that the disk drive needs to be disabled,
5 assigns a part of the disk drives as spares for the disk drive to be disabled; and an anomaly notification step of notifying an occurrence of the disabled state to a predetermined notification destination at a predetermined notification timing; wherein the anomaly
10 notification step may set the notification timing so that the anomaly notification resulting from the sparing processing performed between the disk drives of different kinds is issued earlier than the anomaly notification resulting from the sparing processing
15 performed between the disk drives of the same kind.

Further, this invention may be implemented as a computer program for realizing such a control or as a computer-readable recording media that stores the computer program. The recording media may use a
20 variety of computer-readable media such as flexible discs, CD-ROMs, magneto-optical discs, IC cards, ROM cartridges, punch cards, printed materials printed with bar codes, internal storage devices of computers (RAM and ROM) and external storage devices for computers.

25 Other objects, features and advantages of the invention will become apparent from the following description of the embodiments of the invention taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is an explanatory diagram showing an outline configuration of an information processing system as one embodiment of this invention.

5 Fig. 2 is a perspective view of a disk drive case 200.

Fig. 3 is an explanatory diagram schematically showing an internal construction of the disk drive case 200.

10 Fig. 4 is an explanatory diagram schematically showing an internal construction of a storage device 1000.

Fig. 5 is a flow chart of disk kind management processing.

15 Fig. 6 is an explanatory diagram showing an example configuration of a failure management table.

Fig. 7 is a flow chart of sparing processing.

Fig. 8 is a flow chart of heterogeneous sparing processing.

20 Fig. 9 is a flow chart of failure notification processing.

DESCRIPTION OF THE EMBODIMENTS

Embodiments of this invention will be described in the following order:

25 A. System configuration
B. Disk kind management processing

C. Sparing processing

C1. Failure management table

C2. Sparing processing

C3. Failure notification processing

5 A. System Configuration

Fig. 1 is an explanatory diagram showing an outline configuration of an information processing system as one embodiment. The information processing system has a storage device 1000 connected with host computers HC via a storage area network (SAN). Each computer HC can access the storage device 1000 to implement a variety of information processing. A local area network (LAN) is connected with a management device 10, which may be a general-purpose personal computer with a network communication function and has a management tool 11, i.e., application programs installed in the computer for setting operations of the storage device 1000 and for monitoring the operating state of the storage device 1000.

Installed in a rack of the storage device 1000 are a plurality of disk drive cases 200 and controller cases 300. The disk drive cases 200 each accommodate a number of disk drives (or HDDs) as described later. The disk drives may be 3.5-inch disk drives commonly used in personal computers. The controller cases 300 accommodate controllers for controlling read/write operations on the disk drives. The controller cases 300 can transfer data to and from

the host computers HC via the storage area network SAN and to and from the management device 10 via the local area network LAN. The controller cases 300 and the disk drive cases 200 are interconnected via fibre
5 channel cables (or "ENC cables") on their back.

Though not shown, the storage device rack also accommodate AC/DC power supplies, cooling fan units and a battery unit. The battery unit incorporates a secondary battery that functions as a
10 backup power to supply electricity in the event of power failure.

Fig. 2 is a perspective view of a disk drive case 200. It has a louver 210 attached to the front thereof and an array of disk drives 220 installed
15 therein behind the louver. Each of the disk drives 220 can be removed for replacement by drawing it out forward. At the top of the figure is shown a connection panel arranged at the back of the disk drive case. In this embodiment, the disk drives 220
20 installed in the case 200 are divided into two groups for two ENC units 202, each of which has two input connectors 203 and two output connectors 205. Because two such ENC units 202 are installed in each disk drive case 200, a total of four input connectors 203 and four
25 output connectors 205 corresponding to four paths (also referred to "FC-AL loops") are provided. Each connector has LEDs 204 at an upper part thereof. For simplicity of the drawing, reference number 204 is

shown for only the LEDs of the connector 203[1]. The ENC units 202 may be provided with a LAN connector 206 for a LAN cable and LEDs 207 for indicating a communication status.

5 Fig. 3 schematically illustrates an internal construction of the disk drive case 200. In this embodiment two kinds of disk drives 220 with different interfaces are used. One kind of disk drives 200F has a fibre channel interface (referred to as "FC disk
10 drives") and the other kind of disk drives 220S has a serial interface (referred to as "SATA disk drives"). A circuit configuration that allows for the simultaneous use of different interfaces will be described later. When we refer simply to "disk drives
15 220" they signify disk drives in general without a distinction of an interface. When an interface distinction is made, reference symbols 220F is used for FC disk drives and 220S for SATA disk drives.

The above two kinds of disk drives have the
20 following features. The FC disk drives 220F have dual ports and thus can perform reads and writes from two paths. They also have SES (SCSI Enclosure Service) and ESI (Enclosure Service I/F) functions specified in the SCSI 3 (Small Computer System Interface 3) standard.
25 The SATA disk drives 220S are provided with a single port and do not have SES and ESI functions. It is noted, however, that this embodiment does not exclude the application of SATA disk drives 220S having these

functions.

Shown at the bottom of the figure are side views of the disk drives 220F, 220S. These disk drives have handles 222F, 222S and connectors 221F, 221S for 5 mounting on the disk drive case 200. The connectors 221F, 221S are shifted in vertical position from each other.

As shown at a central part of the figure, the disk drive case 200 has at its back a backboard 230 10 fitted with arrays of connectors 231F, 231S for mounting the disk drives 220. The connectors 231F are for the FC disk drives 220F and the connectors 231S are for the SATA disk drives 220S. The upper and lower connectors 231F, 231S are paired at positions 15 corresponding to the mounting positions of the disk drives 220 and arrayed in a horizontal direction. When the disk drives 220F, 220S are inserted into the disk drive case 200 from the front like a drawer, the connectors 221F, 221S of the disk drives connect to one 20 of the connectors 31F, 231S of the backboard 230 according to their kind. By changing the connectors to which the disk drives 220 connect according to the disk drive kind, it is possible to realize a selective use 25 of circuits that compensate for the interface difference, as described later. The connector difference may also be used for identifying the kind of each disk drive 220. Further, an arrangement may be made so that the kind of disk drive installed is

identifiable from outside. For example, a color of indicator lamp may be changed according to the kind of a disk drive installed or to be installed.

When connected to the connectors, the disk drives 220 are connected to four paths Path0-Path3. In this embodiment, the disk drives 220 connected to Path0, Path3 and the disk drives 220 connected to Path1, Path2 are alternated. This arrangement implements a dual path configuration in which each of the disk drives 220 can be accessed through two of the four paths. The configuration shown in Fig. 3 is just one example, and various other arrangements may be made in terms of the number of paths in the disk drive case 200 and the correspondence between the connectors and the disk drives 220.

Fig. 4 schematically illustrates an internal construction of the storage device 1000. It shows an inner construction of a controller 310 incorporated in controller cases 300 and an inner construction of a disk drive case 200. The controller 310 has a CPU 312 and memories such as RAM and ROM. The controller 310 also has a host I/F 311 as a communication interface with host computers HC and a drive I/F 315 as a communication interface with disk drive cases 200. The host I/F 311 has a communication function conforming to the fibre channel standard, and the drive I/F 315 offers communication functions conforming to the SCSI and fibre channel standards. These interfaces may be

provided for a plurality of ports.

The memories include a cache memory 313 for storing write data and read data written into and read from the disk drives 220 and a flash memory 314 (also called a shared memory) for storing various control software. The controller 310 has circuits for monitoring an AC/DC power status, monitoring states of the disk drives 220, controlling display devices on an indication panel and monitoring temperatures of various parts of the cases. These circuits are not shown.

In this embodiment, two controllers 310[0], 310[1] form the four paths Paths0-Path3 shown in Fig. 3. For the purpose of simplicity, Fig. 4 shows only two loops corresponding to a combination of Paths0 and Path3 or a combination of Path 1 and Path2. These controllers 310[0], 310[1] can switch their paths as shown by dashed lines. For example, the controller 310[0] can access each of the disk drives 220 through either of the two loops, as shown by arrows a, b in the figure. The same also applies to the controller 310[1].

The disk drive case 200 is connected with a plurality of disk drives 220 as described earlier. The FC disk drives 220F are connected to two FC-AL loops through port bypass circuits (PBCs) 251, 252.

The SATA disk drives 220S are connected to two FC-AL loops through a dual port apparatus (DPA) 232, interface connection devices (e.g., SATA master

devices) 233, 234 and PBCs 251, 252. The DPA 232 is a circuit to make each of the SATA disk drives 220S dual-ported. The use of the DPA 232 makes the SATA disk drives 220S accessible from any of the FC-AL loops, as 5 with the FC disk drives 220F.

The interface connection devices 233, 234 are circuits to perform conversion between the serial interface and the fibre channel interface. This conversion includes a conversion between a protocol and 10 commands used to access the SATA disk drives 220S and a SCSI protocol and commands used in the fibre channel.

As described earlier, the FC disk drives 220F have a SES function whereas the SATA disk drives 220S do not. To compensate for this functional difference, 15 the disk drive cases 200 are each provided with case management units 241, 242. The case management units 241, 242 are microcomputers incorporating a CPU, memory and cache memory and collect information on disk kind, address, operating state and others from the disk 20 drives 220 contained in the disk drive case 200. The case management units 241, 242 are connected to two FC-AL loops via PBCs 251, 252 and, according to a SES command from the controller 310, transfers the collected information to the controller 310. In this 25 embodiment, for the controller 310 to be able to retrieve management information in a unified manner regardless of the disk kind, the case management units 241, 242 collect management information not only from

the SATA disk drives 220S but also from the FC disk drives 220F.

The PBC 251 switches the FC-AL loop among three devices connected to the FC-AL loop - the FC disk drive 220F, the interface connection device 233 and the case management unit 241. That is, the PBC 251, according to a command from the controller 310, selects one of the FC disk drive 220F, interface connection device 233 and case management unit 241 and connects it 10 to the FC-AL loop, disconnecting the other two.

Similarly, the PBC 252 switches the FC-AL loop among the three devices connected to the FC-AL loop, i.e., the FC disk drive 220F, interface connection device 234 and case management unit 242.

15 Because of the construction described above, the storage device 1000 of this embodiment has the following features. First, the function of the interface connection devices 233, 234 allows two kinds of disk drives - FC disk drives 220F and SATA disk 20 drives 220S - to be installed in each disk drive case 200. Second, the function of the DPA 232 allows the SATA disk drives 220S to have dual ports. Third, the function of the case management units 241, 242 allows the controller 310 to collect management information 25 also from the SATA disk drives 220S. These features are based on the construction described in connection with Figs. 1-4 and not necessarily essential in this embodiment. In addition to the above-described storage

device 1000, this embodiment can also be applied to storage devices of various constructions including those with a part of the above features excluded.

B. Management Processing by kind of disk

5 Fig. 5 is a flow chart of the management processing by kind of disk to determine the kind of individual disk drives 220, i.e., whether the disk drive of interest is an FC disk drive 220F or a SATA disk drive 220S, and to manage them accordingly. On
10 the left side of the flow chart is shown a sequence of steps executed by the controller 310. On the right side processing executed by the case management units 241, 242 is shown.

When this processing is started, the
15 controller 310 inputs a disk kind check command (step S10). The check command may be issued explicitly by a user operating the controller 310 or management device 10, or an arrangement may be made to take the start of the storage device 1000 as a check command.

According to the check command, the
controller 310 queries the case management units 241,
242 about the kinds of the disk drives 220 installed in each disk drive case 200. Upon receiving this query
(step S20), the case management units 241, 242 identify
25 the kind of each disk drive 220 by checking the connectors to which the individual disk drives 220 are connected. That is, if a disk drive 220 is connected to the connector 231F of Fig. 3, the disk drive is

determined to be an "FC disk drive." If it is connected to the connector 231S, it is recognized as a "SATA disk drive." The case management units 241, 242 notify the check result to the controller 310 (step 5 S24).

The above processing need only be performed by one of the case management units 241, 242 that have received the query from the controller 310. The case management units 241, 242 may also check and store the 10 kinds of disk drives in advance and notify the controller 310 of the check result in response to the query.

Upon receipt of the disk kind check result from the case management units 241, 242, the controller 15 310 stores the check result in a disk kind management table (step S14). The disk kind management table is a table stored in the flash memory of the controller 310 to manage the kinds of individual disk drives 220. A content of the disk kind management table is shown in 20 the flow chart. The disk drives 220 are identified by a combination of a disk drive case 200 number, an ENC unit 202 number and a unique address of each port. For example, a record at the top row of the table indicates that a disk drive 220 at an address "#00" in a disk 25 drive case "#00" and an ENC unit "0" is an "FC disk drive."

The controller 310 repetitively executes the above processing for all disk drive cases (step S16) to

identify the kinds of individual disk drives 220. With the above storage device 1000 of this embodiment, the controller 310 can easily identify and manage the kinds of disk drives even if the FC disk drives 220F and the 5 SATA disk drives 220S are mixedly installed in each disk drive case 200. The controller 310 therefore can take advantage of the features of the FC disk drives 220F and the SATA disk drives 220S in controlling data reads and writes.

10 C. Sparing Processing

The disk kinds of disk drives that have been identified by the methods described above are utilized for the operation and management of the storage device 1000. One example of making use of the disk kind 15 management information on disk drives is sparing. The sparing involves monitoring errors that occur during accesses to individual disk drives, disabling those disk drives which have a sign of impending failure and putting spare disk drives prepared in advance into 20 service before the disk drives become inaccessible.

After the sparing is performed, the controller 310 sends a failure notification to the management device 10 at a predetermined timing in order to prompt the maintenance of the disk drives.

25 For sparing, disk drives stored in the storage device 1000 are grouped into those that are RAID-controlled during normal operation and those that are not used during normal operation but as spares. A

classification between the RAID use and the spare use is stored in a "failure management table" in the flash memory of the controller 310. The failure management table also manages the number of errors in each disk 5 drive and an indication of whether sparing is being performed or not.

C1. Failure Management Table

Fig. 6 shows an example structure of a failure management table. This table records a variety 10 of information about sparing for each disk drive (HDD). Since a plurality of disk drives are installed in each disk drive case (DISK#00-#m) as shown at the top of the figure, the failure management table represents disk drives in a two-dimensional arrangement (with a case 15 number and a serial number in the case). As shown in the figure, disk drives installed in a disk drive case DISK#00 are represented as (0, 0)-(0, n).

Information recorded in the failure management table will be explained. "I/F" refers to a 20 kind of interface of each disk drive, indicating whether the disk drive of interest is an FC disk drive or a SATA disk drive. "Number of failures" means the number of errors that took place during accesses. If this number exceeds 50, it is decided that the disk 25 drive needs sparing. The number "50" is just one example and various other settings may be possible.

"Status" is represented in three states, "normal," "disabled" and "pseudo-disabled." The

"disabled" state means a state in which a disk drive in question is replaced with another disk drive by sparing and removed from service. The "pseudo-disabled" state similarly means a state in which a disk of interest has
5 undergone sparing and is removed from service. The pseudo-disabled state differs from the disabled state in that a failure notification is delayed whereas the disabled state results in an immediate notification of failure. In this embodiment, when the disk drive
10 sparing is performed between the same kinds of interface, this is treated as "pseudo-disabled." When the sparing is performed between different kinds of interface, this is treated as "disabled."

"Sparing" shows a result of sparing performed
15 on a disk drive considered abnormal. "Completed" means that the sparing is completed normally. "Not available" means that sparing cannot be performed because there are no spare disk drives.

In the "spare" column, "yes" indicates that
20 the disk drive can be used as a spare disk drive and "--" indicates that the disk drive is not a spare and is currently used for RAID. Disk drives for which "used as spare" is "ON" are currently in use for sparing.
"Replaced HDD" refers to a disk drive that was found
25 abnormal and replaced with a spare.

In the example shown, since a disk drive (0,
2) has reached the failure number of 50, it undergoes sparing and is replaced with a disk drive (0, 5). The

disk drives (0, 2), (0, 5) are both FC disk drives, so the status of the disk drive (0, 2) is "pseudo-disabled." A disk drive (m, n-1) has reached the failure number of 50 and undergone sparing by which it 5 is replaced with two disk drives (m, n-2), (m, n). Why two disk drives are used will be explained later.

Since this sparing is between different interfaces, the status of the disk drive (m, n-1) is "disabled." A disk drive (0, 4) has reached a failure number of 100 10 but since no spare is available, the sparing field is indicated as "not available."

As described above, the controller 310 executes sparing by monitoring the operating state of each disk drive and using the failure management table. 15 Processing executed by the controller 310 will be explained by referring to a flow chart.

C2. Sparing Processing

Fig. 7 is a flow chart of sparing processing. This processing is executed repetitively by the 20 controller 310 during an operation of the storage device 1000.

In this processing, the controller 310 monitors each disk drive 220 for a sign of possible failure, namely the number of errors that occur during 25 accesses (step S40). When the number of errors exceeds a predetermined value, for example 50, the disk drive 200 of interest is showing a sign of failure and is decided as "having a failure possibility." This

monitoring for a failure possibility is performed for each disk drive.

When a sign of failure is detected, the controller 310 decides that the disk drive in question 5 needs sparing (step S42) and checks if there is any disk drive available for use as a spare (step S44). This check can be made by referring to the failure management table described earlier. It is desired that a RAID group of a plurality of disk drives be made up 10 of those disk drives having the same kind of interface. Thus, when a disk drive fails and needs sparing, it is preferred to check an interface of the RAID group (also called ECC group) to which the failed disk drive belongs. Depending on a result of this check and the 15 kind of spares available, the availability of spares falls into the following three cases:

Case 1: where spares of the same kind as the disk drive with a sign of failure are available;

Case 2: where spares of the same kind are not 20 available but spares of different kinds are available; and

Case 3: No spares are available.

According to the above classification, sparing with a different kind of disk drives is allowed 25 but preceded in priority by the sparing with the same kind of disk drives. In the case 1, the controller 310 selects one of spares of the same kind for sparing (step S46) and updates the content of the failure

management table (step S48). In this case, those disk drives with a sign of failure are "pseudo-disabled."

In the case 2, the controller 310 selects one

of spares of a different kind and performs

5 heterogeneous sparing (step S50). The heterogeneous sparing will be described later in detail because its processing is reverse to and differs from the processing performed when switching from an FC disk drive to SATA disk drive.

10 In the case 3, sparing is not performed but the failure management table is updated (step S48). A disk drive with a sign of failure is assigned a "not available" state in the field of sparing. With the above processing finished, the controller 310 performs

15 failure notification processing according to the result of the finished processing, i.e., notifies the management device 10 of an impending failure (step S60) and exits the sparing processing.

Fig. 8 is a flow chart of heterogeneous
20 sparing processing. This processing corresponds to the step S50 of Fig. 7 and performs sparing between an FC disk drive and a SATA disk drive. When this processing is started, the controller 310 checks the kind of a failed disk drive (step S52). In order to prevent
25 sparing with disk drives having a different interface, a maintenance staff may make an appropriate setting in the failure management table in advance. If such a setting is made, heterogeneous sparing is not performed

when spare disk drives of the same kind are not available.

When an FC disk drive has a sign of failure (step S52), the controller 310 executes sparing by 5 replacing it with a plurality of parallel SATA disk drives (step S54). This sparing is schematically illustrated in Fig. 8. It is assumed that FC disk drives form a RAID with SATA disk drives standing by as spares. When in this condition one of the FC disk 10 drives fails, the controller assigns two SATA disk drives parallelly. Assigning parallelly means storing data distributively in these drives so that the two SATA disk drives are accessed almost parallelly. It is also possible to assign three or more SATA disk drives 15 for one FC disk drive.

Generally, an access speed for SATA disk drives is slower than that for FC disk drives. Thus, by allocating a plurality of SATA disk drives parallelly to one FC disk drive, it is possible to 20 compensate for the access speed difference and minimize a reduction in performance of the storage device 1000 after sparing. Further, the SATA disk drives have lower reliability than the FC disk drives. Therefore, when sparing a FC disk drive with SATA disk drives, the 25 same data on the FC disk drive may be copied to a plurality of SATA disk drives. That is, when sparing an FC disk drive with SATA disk drives, one of the spare SATA disk drives may be mirrored onto the other

spare SATA disk drive.

When a SATA disk drive is failed (step S52), the controller 310 executes sparing by assigning a plurality of FC disk drives serially (step S56). This
5 sparing procedure is schematically illustrated in the figure. It is assumed that SATA disk drives form a RAID with FC disk drives standing by as spares. When in this condition one of the SATA disk drives fails, the controller assigns two FC disk drives serially.
10 Assigning serially means using the second FC disk drive after the first FC disk drive is full. It is also possible to assign three or more FC disk drives to one SATA disk drive.

Generally, the FC disk drives have a smaller
15 disk capacity than the SATA disk drives. Thus, by assigning a plurality of FC disk drives serially to one SATA disk drive, it is possible to compensate for the capacity difference and minimize a reduction in performance of the storage device 1000 after sparing.

20 After executing the heterogeneous sparing in the procedure described above, the controller 310 updates the failure management table according to the result of sparing (step S58) and exits the heterogeneous sparing processing. In this processing
25 the disk drive found to be faulty is "disabled."

C3. Failure Notification Processing

Fig. 9 is a flow chart of failure notification processing. This processing corresponds

to step S60 of Fig. 7, in which the controller 310 controls a timing at which to give a failure notification to the management device 10.

In this processing, the controller 310 checks
5 if there are any "disabled" disk drives (step S61). If a disabled disk drive exists, the controller 310 immediately executes the failure notification (step S67). The disabled state corresponds to a state of a failed disk when sparing is executed between different
10 kinds of disk drives as explained earlier. However, such sparing cannot always compensate well for a performance difference between the different kinds of disk drives even if a plurality of spares are assigned as shown in Fig. 8. Therefore, the controller 310
15 immediately notifies the failure and prompts an execution of maintenance to avoid a performance degradation of the storage device 1000 as much as possible.

When a disabled disk drive does not exist
20 (step S61), the controller 310 then checks for a "pseudo-disabled" disk drive (step S62). If such a disk drive does not exist, the controller 310 decides that there is no need for the failure notification and exits this processing.

25 If a pseudo-disabled disk drive exists (step S62), the controller postpones the failure notification until a predetermined condition is met. As described earlier, the pseudo-disabled state corresponds to a

state of a failed disk drive when sparing is performed between disk drives of the same kind. Since such sparing guarantees the performance of the storage device 1000, delaying the failure notification does not 5 in practice cause any trouble. This embodiment alleviates a load for maintenance by delaying the failure notification under such a circumstance.

If another failure to be notified exists (step S63), it is also notified along with the pseudo-10 disabled drive disk (step S67). The failure notification is also made (step S67) when a predetermined periodical notification timing is reached (step S64). Other timings for the failure notification include a timing at which the number of pseudo-disabled 15 disk drives exceeds a predetermined value Th1 (step S65) and a timing when the number of remaining spares falls below a predetermined value Th2 (step S66). Taking these conditions into account can prevent the failure notification from being delayed excessively 20 after a pseudo-disabled state has occurred.

With the storage device 1000 of this embodiment described above, because sparing between different kinds of disk drives is permitted, an effective use can be made of spares. This in turn can 25 avoid a possible shutdown of the storage device due to a lack of available spares. Since failed disk drives are classified into the disabled and the pseudo-disabled state and the timing at which to issue a

failure notification is controlled according to this failure state classification, it is possible to avoid performance degradation of the storage device 1000 and minimize a maintenance load. After sparing is executed

5 using disk drives of a different kind, a user or maintenance staff, when replacing or adding disk drives, may perform sparing again using the same kind of disk drives as the disabled disk drives. For example, where a RAID group is made up of FC disk

10 drives and a part of the FC disk drives fails and is spared with SATA disk drives, the user or maintenance staff, when replacing the failed (disabled) FC disk drives or adding FC disk drives, may spare the SATA disk drives with the new replacement FC disk drives.

15 This procedure may be performed automatically or manually after the storage device recognizes the replacement or addition of the FC disk drives. Further, if any disk drives are spared with disk drives of a different kind, it is desirable to make this state

20 recognizable on a display or from outside the disk drive case.

A variety of embodiments of this invention has been described above. It is noted, however, that the present invention is not limited to these

25 embodiments and that various modifications may be made without departing from the spirit of the invention. For instance, the circuit for connecting SATA disk drives to the FC-AL and the DPA 32 and SATA master

devices 233, 234 shown in Fig. 4 may be provided on the disk drive case 200 side. While in the embodiments the failure notification is made immediately after a disabled state occurs (step S61 in Fig. 9), this 5 notification timing need not be "immediate" but can be set at any arbitrary timing which is not later than the notification timing of pseudo-disabled states.

It should be further understood by those skilled in the art that although the foregoing 10 description has been made on embodiments of the invention, the invention is not limited thereto and various changes and modifications may be made without departing from the spirit of the invention and the scope of the appended claims.